



Working Paper Series No. 11

On the “Smallest Saleable Patent Practicing Unit” Doctrine: An Economic and Public Policy Analysis

David J. Teece¹ and Edward F. Sherry²

20 January 2016

¹ Thomas Tusher Professor in Global Business, Haas School of Business, University of California at Berkeley, Director of the Tusher Center, Haas School of Business, University of California at Berkeley.

² Chief Economist, Expert Research Associates.

On the “Smallest Saleable Patent Practicing Unit” Doctrine: An Economic and Public Policy Analysis³

I. Introduction

The “smallest saleable patent practicing unit” (SSPPU) doctrine was developed in the context of patent infringement damages awards. It provides that, in calculating patent infringement damages, the damages base should be the *imputed* revenues that the infringer would have earned had all of the actual sales been made of the “smallest saleable patent practicing unit” containing the patented invention. This article does not attempt to summarize the SSPPU doctrine from a legal perspective or examine its legal foundations (or lack thereof). A good summary of the doctrine, which the author characterizes as arising from “a recent series of confusing and contradictory opinions,” and many of the decisions, is found in Sidak (2014).⁴

We show that the doctrine makes no economic sense and is completely at odds with the long standing view that in determining reasonable royalties one should mimic licensing practices in the real world. Such practice usually specify percentages removing royalties on the device revenues or per unit device royalties, but never a SSPPU. The reason is that the value from patented technology is manifested in many places besides components; moreover, transaction cost considerations make SSPPU licenses difficult to monitor and enforce. Accordingly, mandating SSPPU licensing would be a fool’s errand if the goal is a properly functioning natural (or global) system of innovation,

From an economic perspective, the rate and the base should be commensurate with one another. It makes no economic sense to set the rate independently of the base. What matters is the total royalty paid. Many combinations of a rate and a commensurately chosen base yield the same total royalties.

Under the current U.S. patent law regime, the patent holder can seek either lost profit damages (including damages for price erosion) or “reasonable royalty” damages. Under an older patent-law regime, the patent holder could compel the infringer to account for (hand over) all of the infringer’s profits attributable to the patented invention. Even under that regime, it was acknowledged that the patented technology might account for only a portion of the infringer’s profits. It was thought that it would be inappropriate for the patent holder to compel the infringer

³ We thank Peter Grindley for helpful comments. The opinions expressed in this article have been developed over a number of years of our research, both academic and in connection with consulting and litigation projects for a number of clients, notably Intel, Ericsson and Qualcomm. We are solely responsible for the opinions expressed here; they should not be imputed to any other entity.

⁴ Sidak, J.G., “The Proper Royalty Base for Patent Damages,” 10 *J. Compet. Law & Econ.* (2014), 989–1037, at 989.

to account for (and hand over) the entirety of the infringer's profits unless the patented technology was "the basis" for consumer demand for the entire product, so that the infringer would not have earned profits unless it incorporated the patented technology into its products.

Ever since patent law abandoned the idea that the patent holder could require the infringer to account for and hand over all of the infringer's profits attributable to the infringement, and moved to a "reasonable royalty"/lost profits damages framework instead, it has been acknowledged that apportionment is likewise required in the context of "reasonable royalty" damages: no licensor seeks to capture the entirety of the selling price and/or profits associated with the sale of licensed products, but only a fraction thereof.

The issue of royalty/damages base arises in two contexts: licensing and patent infringement damages. We draw a distinction between a *royalty* base (in licenses) and a *damages* base (in patent infringement damages cases), though sometimes the terms are used in a way that blurs the distinction, as when a court talks about a "royalty base" when it really is selecting a damages base for calculating infringement damages. (We are not aware of any examples of the converse usage.)

II. The Smallest Saleable Patent Practicing Unit Doctrine

To our knowledge, the SSPPU doctrine was first enunciated in the *Cornell v. Hewlett-Packard* case in 2006.⁵ The patent in question involved technology for speeding up computation by executing concurrent instructions in a single machine cycle. The technology was implemented in processors, which were built into "CPU bricks," which were in turn incorporated into computers. The patentee's damages expert originally proposed a damages base of the selling price of the entire computer, which totaled some \$32 billion. This was challenged on *Daubert* grounds. The trial court ordered that the damages base be narrowed, but allowed the patentee's damages expert to submit a new opinion. The expert then calculated damages using a damages base of the estimated "CPU brick" revenue (equal to the list prices of the various CPU bricks times the number of each type of brick sold, whether separately or as incorporated into computers), totaling some \$23 billion (a figure the parties did not dispute). The jury verdict based on this testimony in turn was challenged.

In its post-trial decision, the district court concluded that, rather than pursuing a "royalty base claim encompassing a product with significant non-infringing components," the patentee should have based its damages on "the smallest saleable infringing unit with close relation to the claimed invention," which the court said was the processor, not the CPU brick. Some courts have referred to this as the "SSPPU requirement." The trial court found that the appropriate damages

⁵ *Cornell University v. Hewlett-Packard Company*, 609 F.Supp.2d 279 (N.D. N.Y., 2006), available at http://www.oceantomo.com/pdf/Cornell%20UniversityvHewlett-PackardCo_0.pdf.

base was the (imputed) processor revenues, estimated at some \$8 billion (reduced to \$6.69 billion because some of the processors used licensed Intel chips). The need for estimation was due to the fact that “Hewlett-Packard only had actual pricing information for three out of eight infringing processor models. The remaining processor prices were determined using a combination of economic and statistical techniques ...”⁶ As the Federal Circuit said in *VirnetX v. Cisco Systems*, “the requirement that a patentee identify damages associated with the smallest salable patent-practicing unit is simply a step toward meeting the requirement of apportionment.”⁷

We note that, from an economic perspective, the *Cornell* trial court committed a significant conceptual error. The jury was asked to decide both the appropriate royalty rate (which they set at 0.8% as part of the special verdict form) and the total damages (some \$184 million, implying a damages base of \$23 billion, the calculated “CPU brick” figure). The trial court applied the jury’s royalty rate to a much narrower damages base than the jury used, namely the \$6.69 billion in estimated unlicensed processor royalties, saying that the jury’s 0.8% royalty rate was “undisputed” and arguing that applying that rate to the much lower “legally correct” damages base of \$6.69 billion, yielding total damages of \$53.5 million, was “the option ‘most faithful to the jury’s verdict,’”⁸ over Cornell’s objection that the “maximum recovery rule” justified a higher damages award because the jury had heard testimony that a rate of up to 2.5% was appropriate.

The economic fallacy in the court’s ruling is the tacit premise underlying its calculation that the rate and the base can be set independently of one another.⁹ That makes no logical or economic sense; any rate has to be *commensurate* with the base to which it is applied. (By way of analogy, it makes no economic sense to say that “oranges sell for \$1” without specifying what *units* that applies to; is that \$1/apiece? Or \$1/pound? Or \$1/kilo? Or \$1/bag? (and, if so, what size bag?) Or \$1/bushel? In our view, it is extremely implausible that, if the jury *had* been told (as they were not), as a matter of law, that they had to use the imputed processor revenue (about which they presumably heard testimony) as the damages base, they would have chosen the same royalty rate as they chose given the “CPU brick” damages base they actually used. It is much more likely that they would have adjusted the royalty rate upward to compensate (though perhaps not fully) for the use of the smaller base.

Judge Rader’s economic fallacy was to apply the (purportedly) “undisputed” 0.8% royalty rate to a much smaller damages base (of the imputed processor revenues) than the base (of imputed CPU brick revenues) the jury actually used. That exalts legal form over economic

⁶ Id. at *17.

⁷ U.S. Court of Appeals for the Federal Circuit, *VirnetX, Inc. and Science Applications International Corp. v. Cisco Systems, Inc. and Apple Inc.*, No. 2013-1489 (September 16, 2014), available at <http://patentlyo.com/media/2014/09/Virnetx-v-Cisco.pdf>, at *29.

⁸ Cornell, *supra* note 5, at *12.

⁹ The fact that the special verdict form asked the jury to specify both a “reasonable royalty” rate and total damages does not undercut this *economic* point.

substance. The patent holder did not dispute the fact that the jury awarded a 0.8% royalty rate, but the jury did so in the context of a combination of a rate and a base (and the resulting damages award). In our view, the outcome that would have been most consistent with the jury's *overall* award would have acknowledged this.

It is common that the royalty base in a license (on which royalties must be paid) is *not* the same as the damages base used by the courts to award damages (i.e., royalties are often paid on something other than what damages are awarded for). For example, it is common practice in many industries for licenses to call for the licensee to pay royalties on the licensee's worldwide sales of "Licensed Products" (a defined term), regardless of (i.e., without requiring proof of) validity and infringement and regardless of where the product is made, used, or sold. By contrast, U.S. courts can award patent infringement damages only for (a) products "made, used, or sold" in the U.S. that (b) are proven to infringe one or more valid claims of litigated patents. Moreover, patent damages are awarded only during the damages period covered by the suit, which is limited by considerations of statute of limitations and notice, whereas the license may call for royalties to be paid for longer (or shorter) periods. In licenses, the parties can use any royalty base—and structure a deal in any fashion—they feel is appropriate.

In running-royalty licenses/damages awards (as contrasted with lump-sum licenses), one needs to specify the royalty/damages base on which royalties/damages due are calculated. There are two common basic approaches: fixed dollars-per-unit royalties or percentage-based royalties (plus hybrids of these two: e.g., royalties with minimums, per-unit or annual caps, sliding-scales). For dollars-per-unit royalties, the royalty base is the number of units sold. For percentage-based licenses/damages awards, one needs to specify the royalty/damages base to which the percentage is applied. The SSPPU issue arises in the context of percentage-based running-royalty damages awards.

III. SSPPU and the Entire Market Value Rule (EMVR)

Typically, any particular patented technology is responsible for only a portion of the value of the product as sold to end users. The rest of the value is made up of physical inputs, manufacturing and assembling effort, marketing and distribution, entrepreneurial risk-taking, and other patented and unpatented technology. This implies there is a need to *apportion* the selling price and/or profit as between the patented technology at issue and other factors that contribute to value. The need for apportionment has been recognized in patent law for many years and is consistent with how scholars view the sources of profit at the level of the firm.¹⁰ It is not disputed here.

¹⁰ See, e.g., Teece, D., "Profiting from Technological Innovation," 15 *Research Policy* (1986), 285–305; and Teece, D., "Reflections on 'Profiting from Technological Innovation,'" 35 *Research Policy* (2006), 1131–1146.

It would be one thing if the SSPPU were economically “coextensive” with the patented feature, in the sense that either (1) the SSPPU were identical to the patented feature or (2) the patented feature were “the basis” for demand for the SSPPU, but not for the device-as-sold. In such a situation, the argument would be that, even though “apportionment” was necessary for the device-as-sold, apportionment would not be necessary for the SSPPU.

Simply put, if the patented feature is A and the product-as-sold is ABCDE, it is necessary to apportion the selling price of ABCDE to A. If the SSPPU is limited to the patented feature A, then one can look at the selling price of the SSPPU without needing to apportion. (This still leaves the issue of synergies, discussed below.)

However, such a situation is unlikely. More commonly, the SSPPU itself is a complex multi-featured product, the patented feature is only one such feature, and the patented feature is not “the basis” for consumer demand for the entire SSPPU. For example, the SSPPU may be ABC. In such a situation, it is necessary to apportion the selling price of the SSPPU itself as between the patented feature A and other factors (B and C) that contribute to the SSPPU’s selling price.

Put another way, using the SSPPU as the damages base generally does not eliminate the need to apportion, though it *may* make the apportionment task simpler.

Throughout what follows, we will assume that both the SSPPU and the device-as-sold are such complex products.

IV. SSPPU and Hypothetical Licenses

In our view, there is an internal inconsistency between the use of the SSPPU as a damages base and the longstanding view that, in determining “reasonable royalty” patent damages, the court should try to “mimic” what would have been agreed to in a hypothetical license arising from a hypothetical negotiation between the parties, prior to the date of first infringement, on the assumption that the parties agree that one or more of the asserted claims of the asserted patent(s) will be found valid and infringed. This “hypothetical license” or “hypothetical negotiation” framework has long been a part of “reasonable royalty” patent infringement damages law, and is set forth in the fifteenth “factor” in the commonly applied *Georgia-Pacific* analytic framework.

The reason for the inconsistency is that (with one exception, discussed in the next section) we have never seen a license that calls for the licensee to pay percentage-based running royalties based on the selling price of an SSPPU. Instead, the most common practice is for percentage-based running royalties to be calculated based on the price of the product as actually sold by the licensee.

From the perspectives of transaction cost economics and organizational behavior, there are good reasons for such an approach. The parties to licenses want their licenses to be administrable, in the sense that the royalties due are based on data collected by the licensee in its ordinary course of business, and they want the patent holder to have the ability to audit the royalties paid to see if they were calculated correctly. Firms keep their business records based on actual transactions, not on “imputed” revenues calculated as though sales had been made of an SSPPU rather than of the actual product sold.

Mandating that the damages base be the SSPPU not only eliminates virtually all real-world licenses as reference data points for determining a “reasonable royalty,” but also is inconsistent with the way that virtually all actual licenses are structured.

Some U.S. courts have noticed this tension. In *Mondis Technology, Ltd. v. LG Electronics, Inc.*, the trial court said “[i]f [the Entire Market Value Rule] were absolute, then it would put Plaintiff in a tough position because, on one hand, the patented feature does not provide the basis for the customer demand, but on the other hand, the most reliable licenses are based on the entire market value of the licensed products.”¹¹ This point has also been made by other commentators.¹²

We have reviewed hundreds if not thousands of licenses in the course of our scholarly and consulting work. As noted above, it is common that licenses call for royalties to be paid on a royalty base different from the damages base used by courts. We are aware of only *one* real-world license where (a) the license called for a percentage-based royalty, and (b) the royalty base in the license was other than the product as sold by the licensee. (It is an interesting question *why* such licenses are not more common, given the ubiquity of the issue they address, but in our experience, they are extremely rare. We have never seen a reference in the literature to another such real-world license.)

That license was produced subject to a confidentiality order in a litigation context, so we cannot reveal the identity of the parties or the details of the license. But we can illustrate what was involved by means of an example. Consider a licensee that sells products both with and without some *unpatented* feature (e.g., a basic cellphone without a camera versus a cellphone with a camera). The patented feature (e.g., some aspect of cellular technology) is common to both “with and without” products. The license in question called for the licensee to pay a royalty on devices *with* the unpatented feature (e.g., a cellphone with a camera) calculated on a royalty base using the (lower) selling price of an otherwise-identical device *without* the unpatented feature (e.g., a cellphone without the camera feature).

¹¹ Nos. 2:07-CV-565-TJW-CE, 2:08-CV-478-TJW, 2011 WL 2417367, at *e (E.D. Tex. June 14, 2011).

¹² See, e.g., Sidak, G., “The Proper Royalty Basis for Patent Damages,” 10 *J. Compet. Law & Econ.* (2014), 989–1037, at 991, 1020.

The parties to the license acknowledged that a time might come when the licensee no longer sold both products (the “with vs. without” products), and thus there might come a time when the licensee did not have prices available for both products. If that occurred, the license provided that the royalty base would revert to the (higher) price of the product as sold (interestingly, with no change in the royalty *rate* to offset the now-higher base). In such an eventuality, the parties chose not to look at the “closest comparable product” from *another* vendor (e.g., if the licensee stopped selling a cellphone without a camera, the license *could* have provided that the parties would look to another vendor who *did* continue to sell such a product and use that party’s selling price as the royalty base). They chose not to. The advantage of the approach chosen is that the royalty due could be calculated, at all times, solely by looking at the licensee’s product offerings (i.e., to data contained in the licensee’s sales records).

However, we have only seen *one* such license, out of hundreds or thousands reviewed. We have not heard of, or seen references to, any other examples. This suggests that this practice is extremely uncommon. By far the more common practice is to base royalties on the products as actually sold.

V. Psychological Factors: Framing and Anchoring

We suspect that much of the courts’ concerns about the “appropriate” royalty/damages base is driven by psychological considerations, including two well-known features of decision making under uncertainty, known as “framing” and “anchoring.”

“Framing” refers to the fact that the way a problem or situation is presented or “framed” affects individual choices and decisions. For example, people exhibit risk aversion when a problem is framed in terms of losses, but risk preference when a mathematically equivalent (in terms of expected outcomes) situation is framed in terms of gains.¹³

“Anchoring” refers to the experimental fact that individuals tend to make (or have suggested to them) initial “starting points” for their assessments, and then adjust the initial assessment as they receive (or assimilate) additional information. Such adjustments tend to be inadequate, so the starting point tends to affect the final assessment in predictable ways.

Tversky and Kahneman¹⁴ give the example of two groups of high school students, one of which was given five seconds to assess the value of $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$, the other of which was given the same five seconds to assess the value of $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$ (the

¹³ See, e.g., Tversky, A., and Kahneman, D., “The Framing of Decisions and the Psychology of Choice,” 211 *Science*, New Series, No. 4481 (January 30, 1981), 453–458, available at <http://www.uta.edu/faculty/richarme/MARK%205342/Articles/Tversky%2081.pdf>.

¹⁴ See Tversky, A., and Kahneman, D., “Judgment under Uncertainty: Heuristics and Biases,” 185 *Science*, New Series, No. 4157 (September 27, 1974), 1124–1131, available at <http://www.hss.caltech.edu/~camerer/Ec101/JudgementUncertainty.pdf>.

identical product, but presented in reverse order; either can be written as 8! [8 factorial]). Tversky and Kahneman argued,

To rapidly answer such questions, people may perform a few steps of computation and estimate the product by extrapolation or adjustment. Because adjustments are typically insufficient, this procedure should lead to underestimation. Furthermore, because the result of the first few steps of multiplication (performed from left to right) is higher in the descending sequence than in the ascending sequence, the former expression should be adjudged larger than the latter. Both predictions were confirmed. The median estimate for the ascending sequence was 512, while the median estimate for the descending sequence was 2,250 [over four times higher]. The correct answer is 40,320.

Note the wide disparity between the correct answer and *both* of the median estimates.

In patent infringement damages, one “anchor point” is the choice of the damages base. Similarly, whether the issue is “framed” as a (lower) “reasonable royalty” rate applied to a large damages base or a (higher) “reasonable royalty” rate applied to a smaller damages base can make a psychological difference. Mathematically, different combinations of rate and base can yield identical results. There is no mathematical difference between a 1% royalty rate applied to a royalty base of \$10 million, a 2% royalty rate applied to a royalty base of \$5 million, and a 10% royalty rate applied to a royalty base of \$1 million; all three yield the same total royalties, of \$100,000. The rate should be commensurate with the chosen base.

But some courts seem to have been persuaded that, psychologically, juries are more likely to award higher total damages when the issue is framed as a “reasonable royalty” rate applied to a larger base—such as the entire complex product as sold—than when a “reasonable royalty” rate is applied to a smaller base—such as the imputed revenues associated with the SSPPU. Similarly, some courts appear to have concluded that allowing juries to calculate “reasonable royalty” damages using a starting point of a larger damages base—a higher “anchor” point—will yield higher damages than allowing them to calculate damages using a starting point of a smaller (narrower) damages base—a lower “anchor” point. The fact that patent holders often advocate for the use of a larger royalty base, while accused infringers often advocate for the use of a narrower royalty base, lends credence to the idea that *litigants* believe that framing and/or anchoring matters.

As the Federal Circuit put it in *Uniloc*, “The disclosure that a company has made \$19 billion dollars in revenue from an infringing product *cannot help but skew the damages horizon for the jury*, regardless of the contribution of the patented component to this revenue.”¹⁵

¹⁵ *Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1320 (Fed. Cir. 2011) (emphasis added).

Of course, recognizing that “framing matters,” or that anchoring bias exists, provides no independent basis for *choosing* one frame over an alternative frame. If the psychological argument is to be believed—if it applies to collective decision making by jurors, as well as judgmental biases for individuals, a claim for which we are not aware of any empirical evidence—then *any* choice of damages base (*any* choice of “frame” or “anchor”) will “skew the damages horizon for the jury” *relative to* the choice of an alternative base—an alternative “frame” or “anchor.” One still needs to *decide* which approach is appropriate and/or preferable. For example, as noted above, in the *Cornell* case, the choice was between the (actual) revenue at the computer level of some \$32 million, the (imputed) CPU brick revenues of some \$23 million, and the (imputed) processor revenues of some \$8 million (reduced because some processors were licensed). Other than the fact that the first reflects what actually happened, while the second and third are “imputed” figures—a big caveat, especially since the court admitted that five of the eight processor prices had to be “estimated” using statistical techniques—there is no reason to believe that one is inherently superior to another, and in particular no *a priori* reason to believe that one gives “better” or “more accurate” damages than another. Nor is there any *a priori* reason to believe that higher total royalties (which benefit patent holders at the expense of infringers) are worse (or better) from a societal perspective than lower total royalties. That depends on whether one believes that patent holders are over- or under-compensated for their innovations. That cannot be determined by simply talking about the choice of damages base “skew[ing] the damages horizon for the jury.” *Any* one of the choices will “skew the damages horizon for the jury” *as compared to* the other choices. When framing matters, the choice of *which* frame to use must be made *separately* from the mere *acknowledgement* that “framing matters.”

VI. The Value Chain and Patent Exhaustion Issues

In many industries, there is a “value chain” running from suppliers of upstream products through suppliers of downstream products/services to end users. In cellular communications, for example, one value chain runs from suppliers of cellular chipsets to manufacturers of consumer devices such as cellphones and tablets, to cellular service providers, to end users. Another parallel value chain runs from chipset manufacturers to manufacturers of infrastructure devices such as cellular base stations, to cellular service providers, to end users. We talk of cellular service providers as being “downstream” from device manufacturers and chipset manufacturers.

Entities at different levels of the value chain capture different portions of the overall value of being able to use patented technology. That is, there is no such thing as “the value” of being able to use patented technology; value depends on context. Entities at different levels of value chain capture different values (and different fractions of the total value) of being able to use patented technology. The farther “downstream” in the value chain one goes, the greater the cumulative fraction of the total value of using the patented technology that is captured in sales at

that level and “higher” levels in the value chain. Put another way, those “upstream” in the value chain capture only a portion of the total value of being able to use the patented technology.

Patent holders can license their technology at different levels in the value chain. A legal doctrine known as the patent exhaustion doctrine provides (to simplify somewhat) that, once a patent holder has granted an unrestricted license at one level in the value chain, the patent holder cannot extract further royalties from those “downstream” in the value chain from its licensees. Those “downstream” from a licensee are free to use the patented technology in their own products/services without further payment and without needing additional permission from the patent holder. Basically, a license “upstream” “exhausts” the patent holder’s patent rights vis-à-vis those “downstream” from the licensee. If it were not for the patent exhaustion doctrine, a patent holder could enter into multiple licenses with entities at different levels within a single value chain (e.g., with the chipset manufacturer, the cellphone manufacturer, *and* the cellular service provider), extracting some of the value that each entity receives from being able to use the patented technology. The patent exhaustion doctrine prevents this.

This raises several questions. Where in the value chain should patent holders license and/or seek damages? What are the implications of the level in the value chain at which licenses are granted and/or damages sought?

In the cellular communications industry, it is common practice to license at the device level (cellphones and base stations), rather than at either the chipset or cellular service provider levels. Royalties are typically calculated based on the selling prices of the licensed products, rather than as a percentage of the selling price of either chipsets or cellular service.¹⁶ That is, royalties are based on the licensee’s actual revenues, not its “imputed” revenues “as if” devices had been sold at the SSPPU level.

It is conceptually possible to have a license at any level in value chain with a royalty base set at a different level in the value chain (e.g., licensing at the device level, with the royalty base calculated on chipset prices), but we have never seen such a license. The implication is that such an approach to assessing patent infringement damages is inconsistent with industry practice and as such does not match the “hypothetical license” approach. This is further evidence that the SSPPU doctrine is at odds with commercial reality and good public policy.

VII. The SSPPU, Apportionment, and Synergies

Both logically and economically, the rationale for using the SSPPU as the damages base is not mandated by the need to apportion. One could apportion using the selling price of the entire complex device rather than the SSPPU. That said, it may (but need not) be easier to

¹⁶ See, e.g., Stasik, E., “Royalty Rates and Licensing Strategies for Essential Patents on LTE (4G) Telecommunications Standards,” *Les Nouvelles* (September 2010), 114–119.

apportion the price of the SSPPU to some patented technology than it is to apportion the selling price of the device as actually sold to that same technology.

The unstated belief (whether or not it is justified is a separate question) seems to be that apportionment is easier and/or less error-prone—the jury is more likely to “get it right”—if apportionment starts from the selling price of the SSPPU than from the selling price of the entire complex product. The belief seems to be that there is “less to deal with” or “less to strip away” with the SSPPU than with the entire device.

A major concern is that this argument ignores the possible presence of synergies. For example, suppose the task at hand is to value the use of technology A in the context of sales of a complex multi-feature product that has features ABCDE. Suppose that the SSPPU is ABC. Both ABCDE and ABC use A; they differ with respect to *unpatented* features D and E. But the value of being able to use A in the context of ABCDE may be different (due to synergies) from the value of being able to use A in the context of ABC. The fact that, by going to the SSPPU ABC, we have “stripped away” the *unpatented* features DE does not logically or economically imply that it is easier to value A in the context of ABC than it is to value A in the context of ABCDE. We agree that, by doing so, we have stripped certain (allegedly) “irrelevant” features away, but that does not imply that the remaining apportionment tasks are easier and/or less likely to be error-prone.

We are not aware of any empirical evidence for a claim that there is systematic cognitive bias in favor of *excessive* (or insufficient) jury damages awards when using the product-as-sold, rather than the SSPPU, as the damages base, in the sense of damages awards being higher/lower than they “should” be. We are not aware of any independent metric for determining whether damages awards are “excessive.” We are not aware of any evidence cited by courts, or in the literature, to that effect. Absent evidence of a systematic bias, reversing individual damages awards that are seen as excessive does not mandate or justify a change in the general approach to setting patent damages, which the Federal Circuit’s recently adopted SSPPU approach does.

Nor is the “need to apportion” a logical rationale in favor of using the SSPPU rather than using the product-as-sold as the damages base. As noted above, unless the SSPPU product in question is coextensive with the patented technology (highly unlikely), the SSPPU will contain unpatented features as well as the patented one, so one will have to apportion the price of the SSPPU just as one needs to do with the price of the product-as-sold. Using the SSPPU does not eliminate a step. At most, using the SSPPU as the damages base *may* make the apportionment task somewhat easier, as it “strips out” certain unpatented features.

If the concern is the possibility of jury bias against particular types of defendants (e.g., big corporations or foreign defendants), there is no reason to believe that using the SSPPU rather than the product-as-sold would make a difference (the defendant’s identity is the same in either case).

VIII. SSPPU and “Components”

One context in which the SSPPU doctrine has arisen involves cellular telecommunications. Defendants often argue that the SSPPU is the cellular chipset, rather than the cellular handset or base station or the cellular telephone service. Given that (a) chipsets are made and sold at a different (upstream) level in the “value chain” from the cellular devices and the cellular services that also use the patented technology, (b) chipsets typically sell for from one to three orders of magnitude less than cellular devices and cellular service (e.g., a chipset might sell for \$3 to \$10, while a smartphone often sells for \$300 to \$500, and cellular service often sells for \$1,200/year or more), and (c) profit margins for chipsets are often very low, it is not surprising that defendants have argued that the damages base for telecoms patents should be the chipset revenues. One obvious incentive on defendants’ parts is to try to reduce royalty/damages base to a low level. We note in particular that “the value” that chipset manufacturers extract from their use of patented cellular technology is driven by factors such as competition in the chipset market (driven in turn by considerations such as Moore’s law and the very-high-fixed-cost very-low-variable-cost nature of the semiconductor industry generally) and seems almost entirely unrelated to, and is likely to be a tiny fraction of, the value to handset providers and cellular service providers (or end users) of using the same patented technology.

We first note that, to the extent “framing matters,” using the chipset price rather than the handset/base station or cellular service as the damages base, and/or focusing on chipset profit margins rather than handset/base station or cellular service profit margins, tends to “skew” the damages analysis in the infringer’s favor, in ways that are (in our view) not adequately acknowledged by the supporters of the SSPPU approach. Simply put, the value that the chipset supplier puts on being able to use the patented technology is likely to be only a small fraction of either (a) the *cumulative* value, at all levels of the value chain, of being able to use the patented technology rather than the available alternatives; or (b) the value “downstream” from the chipset supplier, at the cellular device/cellular service/end-user level.

As Chief Judge Davis said in his 2014 trial court decision in *Commonwealth Scientific & Research Organization v. Cisco Systems, Inc.*:

... the benefit of the patent lies in the idea, not in the small amount of silicon that happens to be where that idea is physically implemented. ... Basing a royalty solely on chip price is like valuing a copyrighted book based only on the costs of the binding, paper and ink needed to actually produce the physical product. While such a calculation captures the cost of the physical product, it provides no indication of its actual value.¹⁷

¹⁷ No. 6:11-cv-00343, 2014 WL 3805817, at * 11 (E.D. Tex. July 23, 2014).

It would be one thing if one had a reason to believe that the chipset manufacturer had been able to capture most or all of the cumulative value, at all levels of the value chain, of being able to use the patented technology. But that is unlikely to be the case, given the nature of competition at the chipset market, especially if there is “widespread infringement” so that competition prevents firms from building the obligation to pay royalties into the prices they charge for their chipsets.

Similarly, it would be one thing if the patent holder could extract “the value” of being able to use its patented technology at multiple different levels in the value chain, but that is prohibited by the patent exhaustion doctrine. The patent holder gets only “one bite at the apple.” (Admittedly, “downstream” firms are willing to pay higher prices to their “upstream” suppliers if their “upstream” suppliers are licensed, so that the combination of the “upstream” license and the patent exhaustion doctrine protects the “downstream” firm from further liability, than they would be if their “upstream” suppliers are not licensed so that the patent exhaustion doctrine does not apply.)

One second-best approach (in terms of capturing a significant fraction of the overall value of being able to use its technology for the patent holder) is to look as far “down” the value chain as possible. Generally, it is not feasible to license directly to end-users (to capture some of the consumer surplus that they gain from using the patented technology), but it is common to license at the cellular device (handset and base station) level; it is much less common to license at the cellular service level. We are not aware of any licenses that have been granted to cellular device manufacturers that use as the royalty base the price of the cellular chipset.

One clear further disadvantage of the SSPPU approach is that any SSPPU analysis relies on *imputed* revenues and/or *imputed* profits (rather than the infringer’s actual revenues and/or actual profits). In some cases, we do not have actual transaction prices of the SSPPUs. In many contexts, SSPPU sales are infrequent, and are not representative/typical of actual sales generally, so it is not realistic that the parties would have agreed to use the SSPPU price as the royalty base (it is not something kept in the ordinary course of business). The response that “we can estimate the SSPPU price, and assessing damages is not an exact science, so that some uncertainty in the damages base is not dispositive” is not fully responsive, since parties to a hypothetical negotiation would have to come up with some way of agreeing on the royalties due.

IX. Rival-In-Use Goods and Value Cut-Off

In many cases, one good is used as an input into the production of another. In the case of patents, patented technology is used as an input into the production/use of other goods. For (a) rival-in-use, fungible physical goods for which (b) arbitrage is possible, one of the key economic roles of the price mechanism is that it rations goods to their “highest value” uses, so that it is reasonable to infer (excluding varying levels of consumer surplus) that all uses actually made have similar “value.” In such a situation, it is reasonable to assert that there is “a value” of the input, and that any further “downstream” value is not attributable to the input, but is “cut off” by further transformative acts.

By way of analogy, paint and canvas are key inputs into creating a painting, but no one would credibly argue that the value of the painting is attributable to the paint or canvas used. Instead, the value of the paint or canvas is given by the price that the paint or canvas sells for; what the painter chooses to do with the paint, not the paint itself, is what creates the value in the painting. The tube of paint has a market value that is different than the value that the artist can create with the paint, which can vary from artist to artist; differently skilled artists can create differently valued paintings from fungible tubes of paint.

The question is whether this argument carries over to the use of patented technology. We suggest that it cannot. The key difference is twofold: first, different tubes of paint are “fungible” to one another in ways that rival technologies seldom are, and second (and more significantly), different uses are rival, in the sense that a given tube of paint cannot be used to paint multiple paintings; the paint is “used up” in the process of creating the painting. Consequently, the price mechanism is used to “ration” goods to their highest-and-best uses. The possibility of arbitrage eliminates inefficient uses. By contrast, patented technology is not rival-in-use in this sense; “the same” technology” can be used in multiple applications without diminishing its ability to be used in alternative applications. Consequently, there is no need to use the price system to “ration” patented technology across alternative uses, as the alternative uses are not rivalrous. Moreover, as noted above, the technology is not “used up” at the component level; it can continue to provide advantages (relative to alternative technologies) not only at the component level, but also at other “lower” levels of the value chain. The implication is that “the value” of using patented technology can differ across different uses in ways not eliminated by arbitrage or rationing across different uses.

Moreover, for rival-in-use goods, there typically can only be one end user. But multiple firms in the same horizontal “value chain” (e.g., chipset manufacturers, handset manufacturers, cellular service providers, and end users) can infringe a given patent and can each obtain value from their separate infringing uses, while providing products/services to a given end user.

This insight suggests that the obvious economic concern with using the prices of a component-like SSPPU, like the chipset (and especially chipset profit margins), as the baseline for assessing “reasonable royalty” damages is that, at most, one can use the chipset price (and/or chipset profits) to estimate the value *to the chipset manufacturer* of being able to use the patented

technology, *given* competition in the chipset market. But there is no reason to assume (and strong reason to doubt) that chipset manufacturers capture *all* (or even *most*) of “the value” (relative to alternative technologies that could have been used) of being able to use patented technology at other “downstream” levels in value chain

The current value to chipset manufacturers is bounded above by the prices of (and profits on) chipsets, but that does not mean that a “reasonable royalty” should be. U.S. case law acknowledges that a “reasonable royalty” need not be bounded above by the infringer’s actual profits. If current prices/profits are depressed because firms are using patented technology without paying for it (e.g., because of widespread infringement), a “reasonable royalty” can exceed the infringer’s actual profits, possibly significantly.

As an analogy, in pharmaceuticals, while patent protection is in place, a significant fraction of the selling price is attributable to the patent, as is demonstrated when the drug goes off patent and generics enter, when prices fall to something close to production cost (which can be roughly 10% of the branded price).

There is a significant conceptual difference between the value actually received (captured) by an infringing firm and the appropriate value of being able to use patented technology. In the context of widespread infringement, the actual selling price/profit margin may not reflect adequate compensation for infringed patented technology. This is akin to the proposition (recognized in the case law) that a “reasonable royalty” is not capped by the infringer’s profit margin. The price can be raised, if necessary, to provide adequate compensation to the patent holder. The current profit margin reflects current prices and current costs, not what prices *would* be *if* a “reasonable royalty” were paid, which would increase the infringer’s costs to the actual costs plus a “reasonable royalty.”

X. Different Royalties for the “Same Amount” of Use

Another, rarely articulated rationale for the SSPPU approach involves the fact that different products incorporating the patented technology can differ along a number of other dimensions technically *unrelated* to the patented technology, and the prices of different such goods can vary significantly. The argument goes that the patent holder “should not” charge royalties (or collect damages) which vary depending on variations in the selling price of licensed goods, which are caused by the presence or absence of factors that the patent holder did not invent.

For example, Apple sells a variety of iPhones with different amounts of memory. According to its website, Apple currently sells an unlocked iPhone 6s with 16GB memory for \$649, an iPhone 6s with 64GB memory for \$749, and an iPhone 6s with 128GB memory for

\$849.¹⁸ Percentage-based royalties (or damages) based on the selling price of the phone will charge higher royalties (or damages) for phones with more memory (or more unpatented features generally) than phones with less memory (or fewer unpatented features). A holder of a patent on cellular technology charging (say) a 4% royalty would receive \$4 more in royalties (4% of the \$100/unit price difference) for an iPhone 6s with 64GB memory than an otherwise-identical iPhone 6s with 16GB memory, despite the fact that the patent holder’s cellular technology invention has nothing to do technologically with memory and nothing to do with the amount of memory contained in the cellphone. The same is true for many other patented products that sell in versions both “with” and “without” some *unpatented* feature (e.g., a cellphone with and without a camera).

The “argument” goes that, because the patent holder’s invention technologically has nothing to do with the presence or absence of some unpatented feature (such as the amount of memory), the royalties that the patent holder receives “should” not vary with the presence or absence of such features. The (supposed) “corollary” is that charging percentage-based royalties based on the selling price of the units actually sold amounts to charging royalties for the presence of unpatented features, “rather” than for the patented technology itself. Using the SSPPU as the royalty/damages base “strips out” some (but not all) of the variation in price caused by the presence or absence of such unpatented features.

The logical and economic fallacy in this argument is that, in many situations, the *combination* of the patented feature and the presence (or absence) of other unpatented features can result in *synergies* in the *value* that users associate with the product as sold. To take a simple example, adding a camera to a cellphone increases the range of ways that the owner can make use of the cellphone; the owner can now take photos and share them with others over cellular connections in a way that the owner of a camera-less cellphone cannot. This increases the value to the owner of having cellular connectivity. Similarly, the ability to share photos with others over a cellular network enhances the value of the camera functionality, as compared to the value of a stand-alone camera without cellular connectivity that is not capable of such sharing. In other words, even though the two features are *technologically* unrelated, adding the camera functionality enhances the *value* of cellular connectivity, and vice versa.

Similarly, having more memory in a cellphone enables the cellphone to handle larger files in a way that enhances the value of cellular connectivity, and having cellular connectivity likewise enhances the value of having additional memory, because the user can do more with the information contained in the memory if the device is capable of cellular connectivity than if it is not. It is these sorts of value synergies, not merely the technological relationship (or lack thereof) between different features, that is important in assessing “the value” of being able to use patented technology. Put another way, “the value” of being able to use a particular patented technology

¹⁸ Apple Inc., “View Gallery, iPhone Upgrade Program” [webpage] (visited September 30, 2015), available at: <http://www.apple.com/shop/buy-iphone/iphone6s>

can depend on the presence or absence of some other unpatented feature; the two are not *economically* independent of one another, though they may be *technologically* unrelated. This implies that, in such situations, value is not additive, but superadditive. Not surprisingly, in such circumstances a claim that a patent holder “should not” charge royalties that depend on the presence or absence of some unpatented feature “because” the patent holder did not invent the unpatented feature is economically and logically invalid. The “because” claim—that the patent holder did not invent the *unpatented* feature—is of course correct, but the purported conclusion—that the patent holder “should not” charge royalties that vary with the presence or absence of such features—does not follow; the argument is flawed.

We suspect that the source of the fallacy is the (unstated and invalid) premise that there is some such thing as “the value” of being able to use patented technology, independent of the presence or absence of other product features. Or perhaps the source of the fallacy is the (unstated and invalid) premise that value is additive—that “the value” of some complex product is the sum of “the values” of the various features of the product, with no synergies/interactions between those separate “values.”

Why are such premises commonly held? Perhaps because we are used to thinking of pricing and value in the context of physical goods that are rival in use, where prices are used to “ration” goods to different end uses, where there is an established “market price” for the good (so that sellers are unwilling to accept less than the market price and buyers are unwilling to pay more than the market price), and where the price of a basket of different goods just equals the sum of the values of the individual items in the basket (i.e., where prices are additive). But intangible technology is non-rival in use. There is no reason why the use of IP needs to be “rationed” across different users in the way that physical goods need to be “rationed” across rival uses. Using some patented technology in one object does not reduce the ability to use the same patented technology in a different physical object. Charging a higher royalty for the former use than for the latter use does not mean that one of the uses has to be foregone. This is not true for physical goods (such as gold used to plate electrical connections), where using some gold in one physical object (e.g., an iPhone) precludes using that same gold in a different object (e.g., another iPhone). “Arguments” based on (false) analogies between tangible rival-in-use physical goods, where prices ration goods to their highest and best use, and intangible non-rival-in-use patented technology, where such rationing is not needed, are often flawed.

It is easy to assert there are synergies between different features, but it is more difficult to quantify them. Theoretically, it is unlikely that the degree of synergy will be exactly proportionate to the selling price of the product. In an ideal first-best world, one could in theory “titrate” the royalties/damages to the exact degree of synergy. But in practice, one may need to resort to simple, easy-to-administer rules (e.g., percentage-based royalties), relying on the parties to negotiate (or the court to select) a single rate that “smooths out” individuated variations to yield a “blended” rate that is mutually agreeable and (overall, on average) “right.”

XI. On Percentage-Based Royalties

Both the cost and the price (and thus the profit margins) of complex multi-feature products depend on the variety of features included in the product. Thus a price of a complex cellphone with a digital camera, a significant amount of memory, and the ability to wirelessly surf the Internet reflects the variety of features that the cellphone has. The various “macrofeatures” that the consumer perceives typically themselves contain numerous patented technologies, of which consumers are largely unaware, and which are not coextensive with the observable features. (Thus a patent may read on a particular technology contained in the camera function contained in the cellphone.)

A percentage-based royalty rate, or percentage-based “reasonable royalty” damages rate, is applied to the royalty/damages base used to yield the royalties/damages owed. Because the base contains other features besides the patented technology, *any* percentage-based royalties and/or “reasonable royalty” damages award, no matter the level of the royalty rate, is *based on* and thus *reflects* the value of both the patented technology and unpatented features.

Some have argued that using a percentage-based damages approach, whether on the entire as-sold product or on the SSPPU, amounts to “claiming” damages based on the value to consumers (as reflected in transaction prices) of unpatented features contained in the damages base. The suggestion is that royalties/damages “should” only reflect, and/or be based on, the value of the patented technology itself, and not on other features of the product incorporating that technology. This argument is rarely articulated clearly, and its economic basis has to our knowledge not been fully explored. (In our view, the pragmatic difficulties with such an approach are often underestimated and rarely discussed.)

That does not mean that, in asking for percentage-based royalties, the patent holder is seeking overcompensation, as the royalty rate may be sufficiently low that the total royalties being sought/awarded are commensurate with the value of being able to use the patented technology. Given the fact that percentage-based royalties based on actual transaction prices are common in many licenses, for a variety of sound economic reasons, accepting such an argument amounts to abandoning one of the most commonly used approaches to the “hypothetical license” envisioned by *Georgia-Pacific* and may involve abandoning guidance from real-world licenses in the industry.

If one takes the position that percentage-based royalties are not acceptable (for whatever reason), the most likely “solution” is to go with either a lump-sum license, cents-per-unit running royalties, or royalties calculated as a percentage of a price (or profit) differential (discussed further in Section XIX below). Of course, such approaches simply push the problem one step away. The difficult task of setting the lump-sum amount or the per-unit royalty or determining the appropriate share of the price (or profit) differential still remains.

An alternative approach, when the necessary data is available (as it often is not), is to base royalties/damages not on the prices *per se*, but on the price *differentials* between products with and without some feature linked to the patented technology at issue. For example, one might look at the price differential between smartphones with and without camera functionality as a measure of the market value of that functionality, and then further try to apportion that price differential to the patented technology in question. One obvious difficulty is that such “with and without” products may not be offered in the marketplace, especially if the feature becomes ubiquitous (so that consumers expect to have the feature and are not willing to buy products without it). Moreover, that approach simply does not work with features that are common to all products offered.

We are not aware that such an approach has been used in practice; in our review of thousands of licenses, we have only seen one license (discussed in Section VI above) that called for the licensee to pay royalties based on the price of a “basic” product (rather than the more complex product actually sold), and we have never seen a license calling for royalties to be paid on the basis of a price differential. Nor are we familiar with any cases in which such an approach has been adopted by a court.

Still another approach is to base royalties/damages on the *profit* differentials between products with and without some feature (again, this information is often not available). That would more accurately capture the value to the *manufacturer* of offering the feature. The obvious concern from the licensor’s standpoint is that allowing the licensee to pay royalties based on reported profits encourages the licensee to load all sorts of overhead and similar costs onto the profits, so that the reported profits and thus royalties are low. Again, we are not aware of any licenses or any litigated cases that use such an approach.

XII. “Smallest” Unit Issues and “Saleable” vs. “Sold”

What does “smallest” mean in the context of the SSPPU criteria? Looking at the *physically* smallest does not capture what is at issue. Instead, one expects that the relevant criteria is something *economic*, such as “smallest” (lowest) price, cost, profit margin, or value. Starting with a complex device ABCDE containing the patented feature A, there may be multiple choices for “smaller” units, several of which sell for the same price, yield the same per-unit profit to producers, or yield the same value to end users. To illustrate, suppose that products ABD and ABE both sell for the same price (or have the same per-unit profit). “The value” of A in the context of ABD may be different than “the value” of A in the context of ABE. Which should be chosen as “the” SSPPU? To our knowledge, to date courts have not given any guidance as to the criteria they propose to use for selecting “the smallest” saleable patent-practicing unit when there are multiple candidates. In particular, when looking at the value to manufacturers, one presumably wants to look at the profits that manufacturers earn from making

and selling products, whereas when looking at the (unobservable) value to users, one wants to look at the value-in-use to the end user, which is bounded below by the price that the user pays.

Another issue that the courts have not clarified is: does the chosen SSPPU merely need to be “saleable” (whatever that may mean, and whether or not it is actually sold), or must it be *actually sold* separately? Some stripped-down object may be “saleable” in the sense that it could in theory be sold, even if in practice it is not functional (e.g., a car without an engine could in theory be sold, but in practice it is unlikely to find a buyer).

In *Cornell*, for example, as noted above, the district court chose to use the processors as the SSPPU, despite the fact that the defendant did not have actual transaction prices for five of the eight types of processors (so that their selling prices had to be *estimated* using “a combination of economic and statistical techniques”). It was not that they were not “saleable,” but they were not “sold.”

Moreover, even if some item *is* actually sold separately, it may be sold in sufficiently low quantities that using the actual selling price may not be particularly reflective of the value. By way of analogy, automobiles are made up of hundreds or thousands of discrete parts. Automobile manufacturers typically set prices for replacement parts for automobiles, but the sum of the separate prices for the various parts adds up to significantly more than the selling price of the automobile as a whole (and, moreover, adding up the prices of the components ignores the value associated with assembling those parts into a completed car). Many parts are sold separately very infrequently, certainly relative to the volume of sales of assembled cars. In such a situation, “imputing” revenues as though parts actually sold as components of assembled cars had been sold as stand-alone replacement parts will significantly overstate the value to the automobile manufacturer of being able to use the patented technology.

More significantly, the obvious *disadvantage* of merely requiring that the SSPPU be “saleable” rather than actually “sold” is that there may not be any arms-length transaction prices to tie the royalties/damages to. In real-world licensing negotiations, the parties typically want the royalty base to be based on data collected in the licensee’s ordinary course of business, so that the royalties due can be calculated and audited by a third party if disputes about royalty payments arise. Using an unobserved figure as a royalty base merely because the unit in question is “saleable” (but not actually sold) would violate this condition. Using statistical techniques to “estimate” what the selling price would have been is likely to lead to disputes, despite the *Cornell* court’s assertion that “the Federal Circuit allows estimates in damages cases.” We are not aware of any real-world license agreements that call for royalties to be paid on the basis of prices estimated using “a combination of economic and statistical techniques.” To the extent that the court is seeking in setting “reasonable royalty” damages to mimic what would have been agreed to in a hypothetical license, this fact should be considered.

XIII. The “Patent Practicing” Aspect of the SSPPU Criterion

There are yet other unaddressed issues that the SSPPU idea raises. What does “patent practicing” mean? Interpreted literally, it would appear to require that the SSPPU when used as intended satisfies all of the claims of the asserted patents (whether literally or under the doctrine of equivalents). But that does not appear to be how the SSPPU test has been applied in practice. Instead, the courts seem to have substituted some sort of (unspecified) criterion that the proposed SSPPU contain the “inventive element” of the patented invention.

For example, in the *In Re Innovatio* case, the district court used the selling price of the Wi-Fi chipset as the damages base. Yet Claim 1 of Innovatio’s asserted U.S. Patent 5,844,893 is a “method” claim that applies to “a data communication network having a plurality of mobile transceiver units selectively communicative with a plurality of base transceiver units.” A stand-alone chipset is not a “network,” and a chipset does not have “a plurality of mobile transceiver units,” so a Wi-Fi chipset does not satisfy the “network” element of the asserted patent claim. But no one contended that Innovatio’s predecessor-in-interest had invented the concept of a “data communications network” with a “plurality of base transceiver units.” Instead, it seemed to be enough for the district court that the Wi-Fi chipset incorporated the “inventive element” of the Innovatio invention. Alternatively, the court may have concluded that the intended use of the Wi-Fi chipsets was to be built into Wi-Fi-compliant devices that would be used in the context of such a “network,” so that, even if the chipset itself did not infringe, the chipset manufacturer could be liable on contributory infringement and/or inducement to infringe grounds if the chipsets were used in their intended manner.

Some have argued that this is as it should be. The patent claims for a given invention can be written in numerous ways, and infringement (whether literally or by the doctrine of equivalents) depends on the details of the way the claims are written. Some have argued that the return to innovation “should not” depend on the details of the way that the patent claims are written, as that encourages inventors to “game the system” by manipulating the way that claims are written.

Unfortunately, moving away from the precise claims language toward some sort of not-well-articulated “inventive element” test makes it difficult to agree precisely what does and what does not infringe and thus precisely what the SSPPU is.

XIV. The Rise and Fall of the *Lucent* Dicta

In its September 2009 opinion in *Lucent v. Gateway*, a three-judge panel of the Federal Circuit said: “Simply put, *the base used* in a running-royalty calculation *can always be* the value of the entire commercial embodiment, *so long as* the magnitude of the rate is within an

acceptable range (as determined by the evidence).”¹⁹ It went on to say “There is *nothing inherently wrong* with using the market value of the entire product [as the damages base], especially when there is no established market value for the infringing component or feature, *so long as* the multiplier accounts for the proportion of the base represented by the infringing component or feature.”²⁰ This comports with economic common sense: what matters is whether the total royalty is “reasonable,” and different combinations of rate and base yield identical total royalties; the base *per se* does not matter, so long as the rate and base are set *commensurately* with one another, and the rate properly apportions the selling price base as between the patented feature at issue and unpatented features.

Unfortunately, approximately 15 months later, in *Uniloc v. Microsoft*, a different three-judge Federal Circuit panel retreated from this, saying “the passage relied on by Uniloc [the first excerpt quoted above] does not support its position. The Supreme Court and this court’s precedents do not allow consideration of the entire market value of accused products for minor patent improvements simply by asserting a low enough royalty rate,”²¹ citing the *Garretson* case (an accounting-of-the-infringer’s-profits case decided under a much different patent law regime).

Since *Uniloc*, the court has interpreted (one version of) the “entire market value rule” as saying that, in “reasonable royalty” patent infringement damages cases, the damages base *cannot* be the entire selling price of the complex product *unless* the patented feature is “the basis” for consumer demand for the entire product. Such a “negative rule” specifies what is *prohibited* or unacceptable, but provides little guidance as to what is *permitted* when assessing damages. In our view, this version of the Entire Market Value Rule departs significantly from the basic *Garretson* requirement that damages be *apportioned* as between the patented feature and unpatented features of the product unless the patented feature is “the basis” for consumer demand.

XV. The Rise (and Fall?) of the “Smallest Saleable Patent Practicing Unit” Doctrine

The difference between the phrase “patent practicing unit” and the phrase “with a close relation to the claimed invention” has led to some controversy. Does the “smallest saleable unit” have to itself practice the patent (i.e., satisfy all of the patent claims)? Or is it enough if the SSU encompass the key “inventive step,” even if the patent claims specify certain other product or system features that were not themselves part of the “inventive step”?

¹⁹ *Lucent Technologies, Inc. et al. v. Gateway, Inc. et al.*, U.S. Court of Appeals for the Federal Circuit (September 11, 2009), at *61 (emphasis added), available at <http://www.finnegan.com/files/Publication/e16ada36-a5f5-4b40-96ff-01a2f1f1f2bd/Presentation/PublicationAttachment/a1b3bfc3-a4cb-4a7b-86fd-02d8fabc2eac/08-1485%2009-11-2009.pdf>

²⁰ *Id.*, at *62 (emphasis added).

²¹ *Uniloc*, *supra* note 14, at *51.

This issue has arisen in a number of cases involving cellular communications technology. The defendants often argue that the “smallest saleable unit” is the cellular chipset that is the “guts” of the cellphone. But the patent claims often specify aspects (such as a “transceiver” or a “system”) that are not incorporated into the cellular chipset, so that the chipset itself, standing alone, does not infringe the patent (though the chipset manufacturer may be liable for contributory infringement and/or inducement to infringe, when the chipset is used for its intended purpose, by being incorporated into a cellphone that is in turn used as part of a cellular network). The choice makes a big difference to the damages base, because chipsets typically sell for less than \$10, while cellphones can sell for hundreds of dollars, and cellular service can cost a thousand dollars a year or more. Moreover, cellphone manufacturers *buy* chipsets and do not sell them separately (other than incorporated into the cellphone), so the issue of which entity “sells” the product in question sometimes arises.

Many inferred from the discussion in *Cornell* and its progeny that, so long as the patentee had based its damages on the SSPPU—that is, that the proposed damages base was the SSPPU—it would be sufficient, even if the SSPPU itself was a multi-featured complex product and even if the patented feature was not “the basis” for demand for the SSPPU. Thus the VirnetX trial court gave the following jury instruction:

In determining the royalty base, you should not use the value of the entire apparatus or product unless *either* (a) the patented feature creates the basis for the customers’ demand for the product, or the patented feature substantially creates the value of the other component parts of the product, *or* (2) the product in question constitutes the smallest saleable unit containing the patented feature.²²

On appeal, the defendant objected to this instruction, arguing that it “inappropriately creates a second exception that would allow a patentee to rely on the entire market value of a multi-component product so long as that product is the smallest saleable unit containing the patented feature.” The Federal Circuit rejected this jury instruction, saying,

We agree with Apple that the district court’s instruction misstates our law. ... [T]he instruction mistakenly suggests that when the smallest saleable unit is used as the royalty base, there is necessarily no further constraint on the selection of the base. That is wrong. For one thing, the fundamental concern about skewing the damages horizon—of using a base that misleadingly suggests an inappropriate range—does not disappear simply because the smallest saleable unit is used.²³

Where the smallest salable unit is, in fact, a multi-component product containing several non-infringing features with no relation to the patented feature ... the patentee must do more to estimate what portion of the value of that product is

²² VirnetX, *supra* note 7, at *28 (emphasis added).

²³ *Id.*, at *28–29.

attributable to the patented technology. To hold otherwise would permit the entire market value exception to swallow the rule of apportionment.²⁴

Exactly what “more” the patentee needed to do is not specified, but the Federal Circuit seems to be suggesting that the patentee needs to do “more” to reduce the damages base below the SSPPU level (in some fashion), and not merely the royalty rate. Elsewhere, the Federal Circuit said “a patentee’s obligation to apportion damages only to the patented features does not end with the identification of the smallest saleable unit if that unit still contains significant unpatented features”²⁵ and “a patentee must be reasonable (though may be approximate) when seeking to identify a patent-practicing unit, tangible or intangible, with a close relation to the patented feature.”²⁶

The following numerical example may help illustrate our point. Suppose that the infringer’s actual sales are of a “base” unit that sells for \$100/unit, but that the “smallest saleable unit” sells for \$80/unit. Suppose that the base unit costs \$60 to make and sell (for a profit margin of \$40/unit), and the SSPPU costs \$50 to make and sell (for a profit margin of \$30/unit). If the patentee sought a royalty rate of 40% of the selling price of the “base” unit, that would amount to claiming all of the infringer’s profits (not revenues). But if the patent holder seeks (say) a 10% royalty on the base unit, that amounts to a \$10/unit royalty, effectively “apportioning” three-fourths (\$30 of \$40) of the profits on the base unit to unpatented features and only one-fourth (\$10 of \$40) to the patented feature. Similarly, if the patent holder seeks (say) a 12% royalty on the SSPPU’s \$80 selling price, that amounts to a per-unit royalty of \$9.60/unit, effectively apportioning \$9.60/\$30 (or 32%) of the profit on the SSPPU to the patented feature, but leaving the remaining \$20.40 of the \$30/unit profit (or 68%) on the SSPPU (or 25.5% of the \$80/unit selling price of the SSPPU) apportioned to the unpatented features.

The fallacy of the Federal Circuit’s reasoning is the assumption that, if the patentee stops with the SSPPU as the damages base, it is tacitly *apportioning all of the profits* earned on the SSPPU to the patented feature. But that would require the patentee to seek a \$30/unit royalty on the SSPPU, a \$30/80 or 37.5% royalty on the selling price of the SSPPU. Any “reasonable royalty” rate *less* than this *tacitly* apportions only a *portion* of the overall profit on the SSPPU to the patented feature, leaving the rest for the unpatented features. There is no logical reason why the required apportionment needs to take the form of a further reduction (below the price of the SSPPU) in the damages *base*. As the Federal Circuit put it in *Lucent* (in language since disapproved in *Uniloc*), “Simply put, *the base used* in a running royalty calculation *can always be* the value of the entire commercial embodiment, *so long as* the magnitude of the rate is within an acceptable range (as determined by the evidence).”²⁷

²⁴ Id., at *29.

²⁵ Id., at *32.

²⁶ Id.,

²⁷ *Lucent v. Gateway*, *supra* note 18, at *61 (emphasis added).

For example, in critiquing the opinions of the patentee’s damages expert Roy Weinstein, the Federal Circuit in *VirnetX* said, “Weinstein did not even attempt to subtract any other unpatented elements from the [damages] base, which therefore included various features indisputably not claimed by VirnetX, e.g., touchscreen, camera, processor, speaker, and microphone, to name but a few.”²⁸ We do not disagree, but we question why the Federal Circuit felt that it was necessary for him to do so. Whenever the product (whether the “full” product or the SSPPU) contains features other than the patented feature, the selling price presumably reflects the presence of those other features, and thus calculating percentage-based royalties based on the selling price of the product *inherently* involves capturing a portion of the value attributable to such other features. (Conversely, it also involves capturing only a *fraction* of the value attributable to the patented feature.) But mathematically, one can perform the apportionment required under *Garritson* by either adjusting the base *or* adjusting the rate (or some combination of the two). There is no necessary reason why the adjustment has to be made to the sales *base*, and in particular by “subtract[ing]” off some figure attributable to other “unpatented features.” Nor did the Federal Circuit give any useful guidance as to how one would properly go about “subtract[ing] any other unpatented elements from the [damages] base . . .,” as *determining* the amount to be “subtract[ed]” is itself likely to be subject to challenge, especially in contexts where synergies are significant.

In theory, if the infringer sold two otherwise-identical products that differed only with respect to the patented feature, one could use the price *differential* between the two products as a measure of the market value of the patented feature *per se*. For example, if the infringer sold both ABCDE and BCDE, the price differential between the two products could be taken as a measure of the value of the patented feature A. And in such a situation, calculating running-royalty damages based on the price *differential* between the two products would be a more reliable guide to assessing the value of the patented technology (though even here the issue of synergies would have to be dealt with).

In practice, of course, it is rare that an infringer sells two products that differ only with respect to the patented technology. Suppose, for example, that the patent relates to some feature of a digital camera incorporated into a cellphone. It is possible that the infringer sells cellphones with and without cameras, but the patent in question is likely not to be coextensive with the camera feature (as cameras have many features, some patented, others not), so that even if one looks at the price differential between cellphones with and without cameras, it will be necessary to apportion that price differential into those features of the camera covered by the patent and those that are not.

Moreover, unless the seller’s product offerings were sufficiently “rich,” so that it offered numerous products that differed only with respect to a single feature, so that one had the data necessary to calculate a “with vs. without” price differential for each product feature—which in

²⁸ *VirnetX*, *supra* note 7, at *31.

our experience is quite uncommon—there may be no clear way to “subtract” the *value* (as opposed to the *cost*) of other features such as a “touchscreen, camera, processor, speaker, [or] microphone,” so that if the patentee (or its damages expert) *had* sought to “subtract” the value of other “unpatented features,” it is not clear how such an approach could be implemented in practice. Moreover, *pricing* of complex products can be additive, sub-additive, or super-additive with respect to multiple product features. That is, the price of the entire product can just equal, fall short of, or exceed the sum of the individual price differentials associated with each product feature. Generally, there is no economic reason to believe that the selling price of the entire product comprising multiple features is equal to the sum of the separate price differentials for each of the product features separately. Moreover, there may not be a single price differential for any given feature; it may instead depend on the “with vs. without” products being compared. That is, the price differential for a given feature F between product A and product A+F (by which we mean the product containing both A and F) may be different from the price differential between product B and product B+F. If so, which is “the value” of feature F? Compounding the problem is the difference between value and price. For infra-marginal consumers, the value that consumers attach to a product they purchase *exceeds* the price they pay, possibly significantly. Economic theory suggests that, in the presence of synergies, there is no unique “value” of a given feature; value is not simply additive, but depends on context.

In *VirnetX*, the Federal Circuit went on to say,

In reaching this conclusion [that “the patentee must do more to estimate what portion of the value of that product is attributable to the patented technology”], we are cognizant of the difficulty that patentees may face in assigning value to a feature that may not have ever been individually sold. However, we note that we have never required absolute precision in this task; on the contrary, it is well understood that the process may involve some degree of approximation and uncertainty.²⁹

But that argument confuses two different issues: (1) whether it is permissible to use the SSPPU as the damages base if the SSPPU itself is a complex multi-featured product for which the patented feature is not “the basis” of demand (which we will assume in what follows), or whether the patentee must use some other damages base (and, if so, which base? and how should that alternative base be estimated?); and (2) whether “absolute precision” is required in the ultimate task of assessing reasonable royalty patent infringement damages. These are not the same question. The *Cornell* case itself gives an example of the latter: the trial court proposed using the (imputed) processor revenue as the base, despite the fact that H-P did not have processor prices for five of the eight processors that it sold, so that they had to be estimated “using a combination of economic and statistical techniques,” a step that clearly involved “some degree of approximation and uncertainty” and one that would not have been necessary had the

²⁹ *VirnetX*, *supra* note 7, at *29–30.

court used the computer or "CPU brick" damages bases instead. Any reliance on "a combination of economic and statistical techniques" is likely to lack "absolute precision." (The *Cornell* court noted that the parties did not dispute the estimated processor prices, but we suspect that that fact pattern is unlikely to hold in other cases, especially now that the court has set forth the SSPPU rule—which, as noted above, originated in the *Cornell* case.)

The Federal Circuit in *VirnetX* criticized the analysis of VirnetX's damages expert, Roy Weinstein, because "*In calculating the royalty [damages] base, Weinstein ... failed to apportion value between the patented features and the vast number of non-patented features contained in the accused products.*"³⁰ This criticism seems puzzling from economic and public policy perspectives. The task at hand is to determine reasonable royalty damages, which involves (a) determining an appropriate damages base and (b) apportioning the selected damages base as between the patented feature and other features to calculate the reasonable royalty damages. In doing so, it is not logically or economically necessary to do so by "*calculating the royalty base*" by using an "apportionment" approach. We note that, by using the SSPPU rather than the product actually sold, one has already taken a step in this direction. The question is whether one needs to take an *additional* such step. The Federal Circuit seems to be suggesting that the patentee must first start with the SSPPU, then estimate ("calculate") an *apportioned* damages base, and then *further* apportion the already-apportioned damages base (by setting a royalty rate) in determining patent damages. While it is certainly conceptually possible to do so (which is already done to some extent by using the SSPPU as the damages base rather than the product actually sold), imposing a requirement for such a second step increases the complexity of the overall task and runs the risk of introducing an additional possibility of error and an additional step that can be separately challenged. Mathematically, many combinations of a royalty base and a corresponding royalty rate will lead to the same total damages. It makes no logical difference whether one calculates a 2% royalty on a \$50 damages base or a 1% royalty on a \$100 damages base; both yield the same \$1 royalty, though the two approaches may differ psychologically.

We find the Federal Circuit's reasoning in *VirnetX* puzzling. We fully agree that, as the Supreme Court said many years ago in the *Garretson* case (decided under a different patent infringement damages' regime), the patentee "must in every case give evidence tending to separate or apportion the defendant's profits and the patentee's damages between the patented feature and the unpatented features" But mathematically, such an "apportionment" can take many forms. One can start with a large damages base, in which case the apportionment fraction (the percentage of the damages base that should be attributed to the patented feature) will be low; or one can start with a narrower damages base, in which case the apportionment fraction will be higher. There is no logical or economic necessity to start with a royalty base calculated using the SSPPU price (which already removes certain revenues), then apportion that base to remove value attributable to other features (e.g., by subtraction), and finally to apportion the already-apportioned base to focus only on the value of the patented feature. Logically and economically,

³⁰ *VirnetX*, *supra* note 7, at *33 (emphasis added).

one can collapse the two latter steps into a single step without falling afoul of the requirement for apportionment mandated by *Garretson*. Different combinations of royalty rate and royalty base can yield mathematically the same damages amount. There is no *a priori* reason why the task of apportionment is easier using a narrower damages base than a broader one, or why it is either (a) easier or (b) less likely to lead to erroneous results to do it in two stages rather than one.

As of this writing, the latest statement by the Federal Circuit on the SSPPU doctrine came in *CSIRO v. Cisco*.³¹ Cisco had argued for a “rule” which would require starting with the SSPPU in assessing damages. The FedCir rejected that approach:

“The rule Cisco advances—which would require all damages models to begin with the smallest saleable patent-practicing unit—is untenable. It conflicts with our prior approvals of a methodology that values the asserted patent based on comparable licenses.... Such a model begins with rates for comparable licenses and then ‘account[s] for differences in the technologies and economic circumstances of the contracting parties.’... Where the licenses employed are sufficiently comparable, this method is typically reliable because the parties are constrained by the market’s actual valuation of the patent.... Moreover, we held in *Ericsson* that otherwise comparable licenses are not inadmissible solely because they express the royalty rate as a percentage of total revenues, rather than in terms of the smallest saleable unit.... Therefore, adopting Cisco’s position would necessitate exclusion of comparable license valuations that—at least in some cases—may be the most effective way of estimating the asserted patent’s value. Such a holding would often make it impossible for a patentee to resort to license-based evidence.”

“Accordingly, we conclude that the district court did not violate apportionment principles in employing a damages model that took account of the parties’ informal negotiations with respect to the end product.”³²

In other words, the recent FedCir opinion in *CSIRO v. Cisco* acknowledged that many real-world licenses use as the royalty base the product as sold, rather than the SSPPU, and concluded that a rule requiring the SSPPU approach be used would amount to discarding potentially relevant evidence about real-world licensing rates.

In part, the FedCir’s rejection in the *CSIRO* case of the SSPPU “rule” may have been driven in part by the fact that the parties and the court were focused on a cents-per-unit royalty structure, rather than a percentage-based royalty structure. With a cents-per-unit royalty, the royalty base is the number of units sold, and does not depend on the selling price, whether of the

³¹ *Commonwealth Scientific and Industrial Research Organization v. Cisco Systems, Inc.*, U.S. Court of Appeals for the Federal Circuit, No. 2015-106 (December 3, 2015), available at <http://www.cafc.uscourts.gov/sites/default/files/opinions-orders/15-1066.Opinion.12-1-2015.1.PDF>

³² *Id.*, pp. *14–15 [internal quotations and citations omitted].

product-as-sold or the SSPPU. With such a royalty structure, there is no “apportionment” of the licensee’s revenues needed.³³ That response begs the question however, as it sidesteps the general requirement of apportionment.

It remains to be seen how the FedCir’s *CSIRO* decision will affect future use of the SSPPU approach. But we anticipate that the FedCir’s explicit rejection of a proposed “rule” requiring its use, and the explicit endorsement of both a license-based approach (coupled with the fact that very few licenses use the SSPPU as the royalty base) and a cents-per-unit royalty structure will dramatically cut back on its use on a going-forward basis.

XVI. Pragmatic Uses:

It would be one thing if, in valuing the use of patented technology A in the context of the sale of a complex device ABCDE selling for P, we had as a comparison reference point the price of an otherwise-identical device that did *not* incorporate the patented technology BCDE selling for a lower price P’. One might then try to calculate “the value” *to the end user* of being able to use A in the context of ABCDE as the price *differential* P–P’ between the patented product ABCDE and the otherwise-equivalent *unpatented* product BCDE. Similarly, when looking at “the value” *to the manufacturer* of being able to use A, one might want to look at the *profit differential* $\pi - \pi'$ between the per-unit profit π that the manufacturer earns selling ABCDE and the per-unit profit π' that it earns selling BCDE.

Of course, BCDE is not an SSPPU for technology A, because it does not practice the patented invention. Nor is it common that firms sell (or that consumers have available) such “with and without the patented feature” products. (There is a greater chance that firms might sell two products (e.g., ABCDE and BDE), which differ with respect to *both* the patented feature A and some unpatented feature C.)

We are not aware of any license that calculates royalties as a portion of a price (or profit) differential. Nor are we aware of any cases that calculate patent infringement damages based on a price (or profit) differential, despite language in some cases that such an approach is appropriate.

Another complication comes from the internal schizophrenia associated with the inconsistency between the “hypothetical negotiation” framework and legal constraints on damages. If one took the concept of a “hypothetical license” seriously, one would pay attention to such real-world licensing considerations as administrability. The parties to a license want the ability to base royalties due (and the licensor typically wants the ability to audit royalties to make

³³ As the FedCir put it, the SSPPU principle “is inapplicable here, however, as the district court did not apportion from a royalty base at all.” *Id.* at *13.

sure that the proper amount of royalties have been paid) on information collected and kept by the licensee in the ordinary course of business, which typically requires that royalties be based on actual sales, not on what the sales “would” have been had they been made (as they were not) at the SSPPU level.

To take one common example, patent infringement damages are available in the United States only for products “made, used, or sold” in the U.S., despite the fact that pragmatic considerations often cause the parties to agree for the licensee to pay royalties based on its worldwide sales, in part because it is often difficult for the parties to know whether products made and originally sold overseas will end up being “used” in the United States. For example, the licensee may make and sell DRAMs overseas, but some of those DRAMs will end up being incorporated by the licensees’ customers in computers which are in turn sold in the U.S. The licensee’s own sales records typically do not have the information needed to identify the location of the final sale or use. As such, the parties may agree that the licensee will pay royalties based on its worldwide sales, despite the fact that the licensor concededly does not have patents in all countries in the world. Such provisions are often explicitly stated as being “for the convenience of the parties.” The royalty rate typically is adjusted (downward) to reflect the “mismatch” between the less-than-worldwide patent coverage and the worldwide royalty base.

Or the parties may agree that royalties will be paid on all “Licensed Products” (a defined term) without requiring proof that all such products infringe one or more valid claims of the licensor’s patents, whereas the courts will only award patent infringement damages for products shown to infringe one or more valid claims of the Patents in Suit. After all, one purpose of a license is to give “patent peace,” and in a world in which the patent holder’s patent portfolio and the licensee’s product mix are both changing over time, it would be costly, time consuming and disruptive to make the obligation to pay royalties depend on proof that particular products infringe valid claims (as the parties may well dispute validity and infringement on a product-by-product basis).

Similarly, parties to licenses want their licenses to be “administrable” based on verifiable information kept in the ordinary course of business. One disadvantage of the SSPPU approach to damages is that it does not comport with what parties to a license would agree to. Licensees do not keep business records “as though” they had sold the SSPPU; they keep records of what they actually sold. And a licensor that wants to conduct a royalty audit to see if the appropriate amount of royalties was paid wants to rely on information as kept in the licensee’s ordinary course of business.

As noted above, the *Cornell* court had to rely on “economic and statistical techniques” to *estimate* the selling prices of *five of the eight* processors at issue. It had to do so only because of its decision to use the processor price as the damages base; no such estimation would have been necessary if the court had used the actual transaction prices of the computers as actually sold, which were reflected in the defendant’s business records. In such situations, reducing the

“framing bias” comes only at a cost in accuracy (namely, using estimated figures rather than actual transaction records). The courts have not explained why they believe the former outweighs (or should outweigh) the latter. Nor have they clearly articulated how they evaluate the *tradeoffs* associated with choosing one damages base over another, nor the criteria they use in making such decisions.

XVII. The SSPPU, FRAND Licensing, and the New IEEE Intellectual Property Rights Policy

Like many standards-setting organizations (SSOs), the Institute of Electrical and Electronics Engineers (IEEE) has an Intellectual Property Rights (IPR) policy that asks holders of patents that may be “essential” to practice an IEEE standard to commit to making licenses available on “reasonable and non-discriminatory” (RAND) terms. (If such a RAND commitment is not forthcoming, the IEEE will elect not to incorporate the patented technology into the standard.) Originally, the IEEE provided little in the way of guidance in terms of what constituted RAND terms. However, in February 2015, the IEEE-SA adopted a number of controversial changes to its IPR Policy, over the objections of a number of major patent holders.³⁴

Among the changes was a provision that a “reasonable royalty” should be interpreted in terms of “the value that the functionality of the claimed invention or inventive feature within the Essential Patent Claim contributes to the value of the relevant functionality of the *smallest saleable Compliant Implementation* [of the standard] *that practices the Essential Patent Claim,*” coupled with the assertion that a “Compliant Implementation” can be a “component” or “sub-assembly” that practices the standard.³⁵ The new IEEE-SA IPR Policy further provides that RAND should consider “the value that the Essential Patent Claim contributes to the smallest saleable Compliant Implementation that practices that claim, *in light of the value* contributed by *all* Essential Patent Claims for the same IEEE Standard practiced in that Compliant Implementation.”³⁶ To our knowledge, no other SSO has to date adopted such provisions into its IPR policies.

Concerns have been expressed that (a) the effect of various aspects of the new policy will be to reduce the royalties that implementers will have to pay for their use of others’ essential patents, thereby reducing the returns to successful patented innovation incorporated into IEEE standards; and (b) firms with particularly valuable patented technology may elect not to make RAND commitments under the new policy.

³⁴ See, e.g., Crouch, D., “IEEE Amends its Patent (FRAND) Policy,” Patently-O, guest post by Professor Jorge L. Contreras (February 9, 2015), available at: <http://patentlyo.com/patent/2015/02/amends-patent-policy.html>.

³⁵ IEEE Standards Association, “IEEE-SA Standards Board Bylaws,” Sec. 6.1, under “Reasonable Rate” (emphasis added), available at <http://standards.ieee.org/develop/policies/bylaws/sect6-7.html>

³⁶ *Id.*

It remains to be seen how the new IEEE's policy's focus on the "smallest saleable Compliant Implementation" will play out in practice.³⁷ To our knowledge, the issue has not yet arisen in practice, nor has it been litigated.

Conclusion

The "smallest saleable patent practicing unit" (SSPPU) doctrine is a recently developed court-made doctrine applied in assessing patent damages in cases where the product-as-sold is a complex multi-featured product for which the patented technology at issue is not "the basis" for consumer demand. It makes no economic sense and is at odds with licensing practices. It appears to be driven by concerns about biases potentially caused by framing and/or anchoring by juries. An even bigger concern is that it is being used to surreptitiously lower reasonable royalty damages, and hurt the innovation process.

We have discussed a number of its limitations, including the fact that very few real-world licenses comport with the SSPPU doctrine, making it difficult to appeal to the terms of real-world licenses in assessing reasonable royalties. We note that the SSPPU doctrine presents conceptual problems when applied to components produced at one level in the value chain. We also note that the SSPPU doctrine typically does not eliminate any step in the requirement that damages be apportioned to the patented technology. One standards-setting organization, the Institute of Electrical and Electronics Engineers, recently adopted the SSPPU as part of its Intellectual Property Rights policy. It remains to be seen how the doctrine will play out in practice. We predict that it will have a significant deleterious impact on innovation. That's to be expected with any judicial ruling that doesn't make economic sense.

³⁷ The IEEE-SA also announced: "While the IEEE-SA Patent Policy *recommends* considerations for use in determining a Reasonable Rate, these considerations are *not mandatory*" (emphasis added) (IEEE-SA, "Understanding Patent Issues During IEEE Standards Development" (September 1, 2015), available at: <http://standards.ieee.org/faqs/patents.pdf>). Given this caveat, it is not clear how the new "should" language will be interpreted (by licensors, licensees, or the courts) in practice on a going-forward basis. One possibility is that, given the "not mandatory" language, patent holders will seek to clarify their FRAND declarations to make their intentions clearer.